Collaborative Browsing in the World Wide Web

Proceedings of the 8th Joint European Networking Conference, Edinburgh, May 12.-15. 1997

M.S.E.E. Gabriel Sidler, ETH Zurich
Dr. Andrew Scott, Lancaster University
Dipl. Physiker Heiner Wolf, University of Ulm

Table of Contents

Abstract

I. Why Collaboration on the WWW?

II. The CoBrow Approach
- II.A. Logical Vicinity
- II.B. Meeting Places

III. Applications
- III.A. Prototypes
  - III.A.1 Lightweight Conferencing
  - III.A.2 Web only Approach
  - III.A.3 Java Approach

IV. Conclusions

V. Ongoing and Future Work

VI. Acknowledgments

VII. References

Author Information

Abstract

The World Wide Web (WWW) is today the most successful service of the Internet. The richness of information available combined with easy access to this information makes it a premier information gathering tool for researchers and consumers, however, the model of today's WWW does not include the users. The WWW is a purely information focused environment, consisting of documents and links between these documents. The virtual world formed by the linked information on the WWW is completely separated from the world of its users. The CoBrow project [6] proposes to extend the model of the WWW to include its users. This will enable many new applications like WWW based conferencing, help desks, online presentations, online tours, and group entertainment. We believe that the WWW is well suited to become the unifying platform for synchronous, interactive collaboration across the Internet.

CoBrow has developed a model of a WWW that includes its users. The approach chosen is to associate users of the WWW with a location in the WWW; this association is based on the document that the user is viewing. Furthermore, users can have attributes like Interest, Language, Time of Presence etc. Based on the location of users in the WWW and these attributes virtual neighbourhoods

are formed. Users within the same neighbourhood are made visible and can collaborate with each other through conferencing and application sharing tools.

Three prototypes were implemented and allowed us to gain experience with different approaches and technologies. An important aspect is the seamless integration into the current WWW infrastructure without requiring changes to existing protocols.

I. Why Collaboration on the WWW?

The WWW was originally envisioned by its designers to become a collaborative tool for the Internet, however, the early implementations were little more than a tool to publish and retrieve documents to and from the Internet. Nevertheless, from the users' point of view it was a big step forward compared to the tools that existed at that time for the search and transfer of documents. The graphical user interface that could be operated by novice users, the capability to link documents and the platform independence of the content format were the main advantages.

Collaboration on the Internet is certainly not a new research topic. Initial efforts focused on asynchronous collaboration and produced such successful tools as Usenet News, Email, mailing lists, and the file transfer protocol. More recently the focus of research has shifted to synchronous collaborative tools such as Internet telephony, audio-, and video conferencing tools. It was proven that the Internet can support such real-time collaborative services if sufficient resources are available.

Most of today's synchronous collaborative tools are not integrated into the WWW. Integration is to be understood in two ways. First, the tools are typically standalone applications outside of the WWW client software. More importantly, there does not exist a mechanism to become aware of other users on the WWW and establish contact with them.

During the past few years the WWW has experienced enormous growth [3]. The number of WWW servers providing information is now at 275,000 [1] and still growing exponentially. Today the WWW provides an extremely rich virtual world of linked information, however, WWW users are not a part of this virtual world. In today's model of the WWW users are watching this exciting and colorful world of information from the sidelines. They request information by specifying a URL and then their WWW client fetches the requested document from the WWW and displays it. This separation of the information world from the world of the information consumers and producers clearly limits the usefulness of the WWW. WWW users are not aware of each other and therefore have no means to communicate. CoBrow attempts to unite the two worlds by developing a model of the WWW that includes its users. This will enable new collaborative services including WWW based conferencing, help desks, presentations, online-tours, group entertainment, etc. We believe the WWW is well suited to becoming the unifying platform for synchronous, interactive collaboration across the Internet for several reasons:

- **Multiplatform Support:** Designed to run in a heterogeneous environment, the WWW supports a large variety of hardware, software, and networks. This ensures interoperability of tools, content, and services.
- **Universal Coverage:** Access to the WWW is virtually universally available. A phone line is all that is needed.
- **Publicly Available:** The WWW is owned by its users. Access is unrestricted and the only costs to participate are connectivity and operational costs.
- **Large Audience**: The WWW is the most successful online service. The number of users continues to grow exponentially.

These four features make the WWW a unique environment, capable of becoming the universal collaborative platform of the future.

II. The CoBrow Approach

The WWW is a meshed network of hypermedia documents. WWW users have random access to virtually every document from their desktops. The WWW is ideally suited for browsing through the information space, but many people are searching for specific, possibly related, information. They often start their search at starting points that are known to be linked to relevant information.

Alternatively people can submit keywords to search engines and hope to get a significant answer directing them to the desired information. There is a vast amount of information available on the WWW, but information about any given subject is not concentrated on one server, it is distributed over the Internet. Chunks of information are available from many different sources. WWW users who search for information usually get access to a small subset of the information resources available. One of the most important reasons for this is that people do not have enough time to explore all the resources offered by search engines. Everyone searching for a certain piece of information will explore a different subset of resources. Over time everyone learns about useful information sources, but the knowledge gained covers only a subset of that available, however, the combined knowledge of all the people interested in the same subject covers virtually all the information available at all sites.

Searching for information would be much more efficient if people were able to share their knowledge about information resources. The problem is to bring together people who can benefit from each other. The WWW is crowded by several million people world wide, and somewhere in the order of one million people are estimated to be active concurrently.

If someone browses for information, there is a high probability that someone else is interested in the same subject at the same time, but people browsing the WWW are unaware of the presence of any fellow browsers. Even if they are submitting the same keywords to a search engine or are searching at the same starting point, there is no possibility of getting in touch with each other. It is the goal of CoBrow to bring these people together. CoBrow is a distributed tool set providing support for collaborative browsing in the WWW. It essentially generates lists of people (clients) who share the same interests, can learn from each other, or are interested in meeting each other. The decisions made to generate these lists are based on information available about the users.

II.A. Logical Vicinity

To overcome the web surfer's loneliness and bring him together with other users it is necessary to define some criteria for the grouping process. Many people having different knowledge, interests and intentions are browsing in the WWW. Hence, decisions have to be made concerning the grouping of those who might benefit from each other. Different evaluation criteria (different types of information about users) are combined to form a virtual space (vicinity space). A *vicinity metric* defines the distance of users in this space. The vicinity metric is a combination of four simpler metrics:

- space metric
- semantic metric
- time metric
- metric of user interests

These four metrics can be represented as orthogonal axes in the vicinity space (Figure 1). Despite their orthogonality, only combinations of metrics are adequate to limit the number of potential participants in a group.

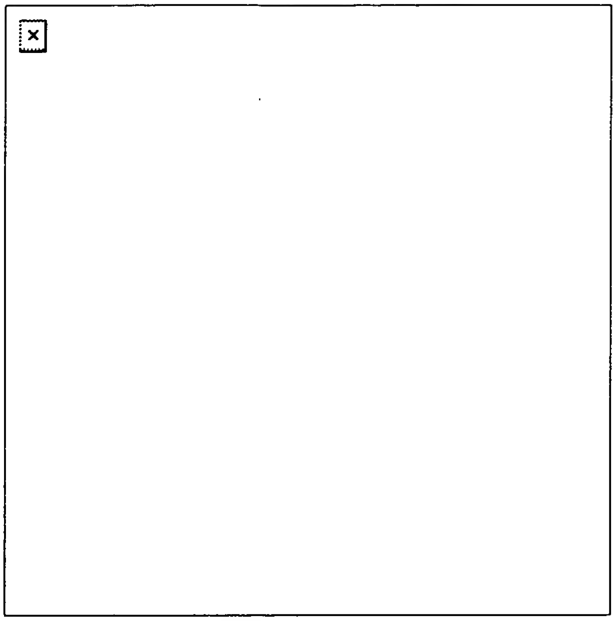

Figure 1: Metrics form a 4-dimensional space

Space Metric: The document structure in the Web can be represented by directed graphs with the HMTL documents as vertices and the links between documents as edges. The distance is calculated as the smallest number of hypertext references that link documents. Using the space metric enables groups to be formed of people browsing within a spatial environment defined by the links between documents. The application of the space metric combined with a time metric could, for example, serve as the basis for an online-help system in the Web.

Reliable tracking of user movements through the Web is crucial to detecting a spatial vicinity among users. CoBrow needs up-to-date information about the document being displayed by the user's WWW client, a task complicated by the number of web servers a user may contact and by the lack of any proper record of a user's state at a server. A number of ways of addressing the problems of maintaining persistent state information have been tried within the project and are discussed in the next chapter.

Semantic Metric: The semantic metric is based on document semantics. The semantics can be covered by introducing weighted hyperlinks or by using content-based information from the documents. Combining the semantic metric with spatial metrics considerably reduces search spaces.

For the weighted link metric, hyperlinks between HTML documents are tagged with weight attributes. These weight attributes expresse the intensity of the relationship between documents as a hint for content correlation. A high weight implies a strong correlation. A threshold value for the computation of the logical vicinity brings together only those users whose set of visited documents fulfil strong correlation criteria. Link weights can be added to documents by the author or by analysis tools based on, for example, keyword matching algorithms.

The document semantic-based metric is more challenging to bring about. Documents which are not necessarily linked, but deal with the same topic, should be included in a vicinity. To obtain the relevant content of a document, information retrieval algorithms are necessary. The algorithms required are very similar to those used by Internet search engines. A simple CoBrow could use the services of search engines or just rely on keywords contained in HTML META-tags.

Time Metric: People browsing at the same time or in overlapping intervals of time are in a temporal vicinity. The time metric is a supporting metric for the other metrics discussed and is not particularly useful on its own. Using time criteria alone would bring more or less all concurrent users of the WWW together in one vicinity.

Temporal and spatial information about users browsing the WWW is available and we are currently exploring a time metric which takes into account the length of time documents are displayed at the client as well as the time at which presentation took place. The longer a user reads a document, the longer it is associated with the user, and subject to spatial-user matching. We are considering display duration as an indication of importance. We are aware that a long duration may also be an indication of complexity, or simply of a break for lunch.

Figure 2: Value of time metric is at a maximum during presentation of a page to the user and decreases thereafter

User Interest Metric: Whereas the previous metrics have been WWW server and document-centred, the user interest metric is human-centred. Here, vicinity is defined to correlate users according to attributes assigned to them. Attribute matching could be based on shared interest, the ability to speak a common language, membership of the same cultural group, and many other definable characteristics. Attributes are collected at the client side and exchanged with the CoBrow system automatically, or on request from the user.

The creation of user interest profiles is beyond the scope of the core CoBrow service. We are exploring mechanisms to create user interest profiles automatically using information available at the client side, for example, the news reader configuration or statistics about WWW traffic. If the CoBrow service is provided in connection with the service of a search engine then CoBrow can directly profit from the information available from WWW users. Users submit keywords to search engines in order to ask for information about a certain topic, and by submitting keywords users implicitly express their current interests. The normal search engine will use these keywords to find related documents and CoBrow, as a search engine for people, will work on exactly the same keywords to find people with related interests. A user who enters a discussion forum on the WWW also expresses his current interests by joining. The subject of the forum entered clearly indicates his interest. In this case an interest based metric is applied very easily. The list of users generated by CoBrow is just the set of users in the forum.

II.B. Meeting Places

We use *meeting places* as a metaphor to represent the vicinity. Meeting places have a different appearance depending on the metrics used for their creation. When using space or semantic metrics in conjunction with time, meeting places appear spontaneously on the client's browser interface as people move through the web. The meeting place contains information about other users in the vicinity and offers synchronous collaboration tools for establishing contact with these users. If user interest is adopted as a metric for a meeting place, forums where people can meet will instantly appear when users with similar interests are in close proximity.

III. Applications

There are many situations where discussion based on web pages would be useful, for example:

- In a research environment it would be useful to be able to discuss ideas and theories with researchers in similar fields, this can be done by prior arrangement with existing conferencing systems, but how often is important work being done by similar groups missed because researchers in a related field are forgotten about or never even considered? The web can help here as it is far easier to discover related work via the web than, for example, through Usenet news. Discovery of related topics and people is becoming easier in the web with the development of search engines that regularly trawl through the Internet constructing lists of related web pages, admittedly these are still fairly unsophisticated but development of these systems has only really just begun;
- In a commercial environment it would be nice to have support for user queries. A number of companies have email and form based feedback systems but simple chat tools or better still multimedia conferencing support would be ideal for user support. Users could be guided through user interfaces, shown new products in greater detail than possible with still images and simple sales videos, etc.
- Also the numerous cases where users would simply like to have a chat about something they have seen on a web page. Perhaps because they are interested in finding out more about the subject or they would like to let the page author know about something they have been doing in the same area.

In other words a synchronous multimedia conferencing system built around the web would be useful in providing communication for the mutual benefit of users with similar interests, particularly when these users were not previously aware of each others existence or of their shared interest.

III.A. Prototypes

Due to the variety of application domains in which a CoBrow system could be used, and the number of alternative approaches to the implementation of such a system, each partner was given the task of building a simple prototype using one approach. Three experimental prototypes have been implemented and our experiences and the lessons learned during the development of these systems are now driving the design of a fully functional CoBrow system.

A secondary, but no less important, reason for developing a range of early prototypes was to enable members of the project team to get feedback from real users. Many of the issues faced during the design of our final system depend on how real users respond to different styles of interfaces and we were particularly interested in discovering what people's reactions would be to being tracked, presented with the names of users with similar interests, etc.

III.A.1 Lightweight Conferencing

An issue that came up quite early in discussions about a CoBrow was the amount of pre-installed software on client machines that would be necessary to support collaborative browsing. In an ideal world the whole system would be entirely self contained and everything would appear within a web browser. Like most ideals this is difficult to realize with any degree of success, however, today's web browsers offer a number of features which can be used to go some way towards a lightweight conferencing system with little pre-installed software.

The lightweight conferencing prototype makes use of animated-GIF functionality to provide video images on standard web pages. Early trials of this approach to adding video took place at the University of Ulm and resulted in the model railroad demonstration [5].

A CoBrow system using animated-GIFs, or *web-video,* in this way depends on each client host running

a small daemon application that provides an interface to the world that looks like a WWW server (Figure 3). This web-video server simply distributes a page containing an infinitely long animated-GIF consisting of frames taken from the host machine's video camera.

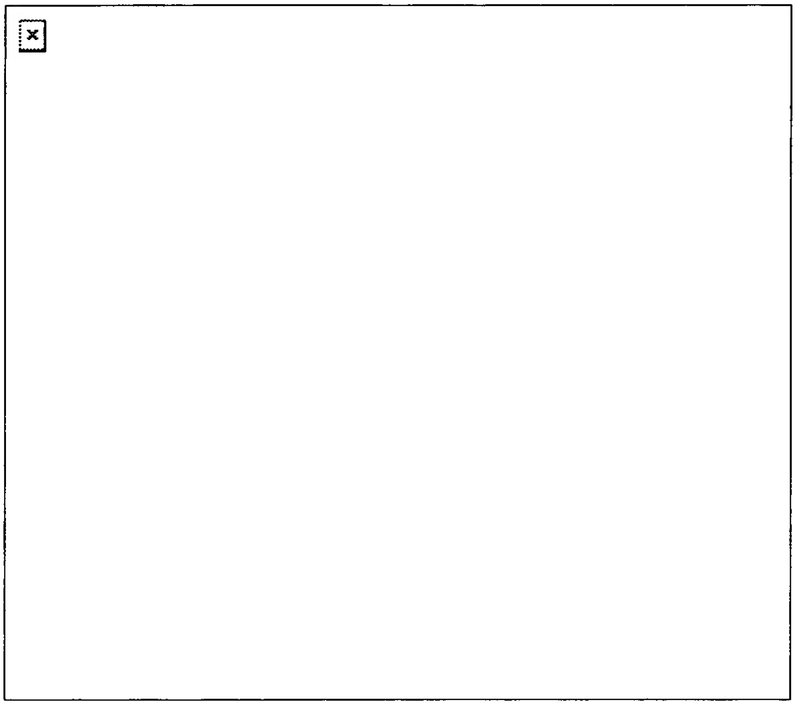

Figure 3: Lightweight conferencing using personal web-video servers

A problem with this approach is that a separate web page must be displayed for each user in a meeting place and while good for small numbers of users, larger groups require too much desktop as the number of browser windows increases.

We are currently using a Java based mechanism to get information about the start and end of the presentation of documents. A small Java applet is inserted into every HTML document subject to the space metric on CoBrow enabled servers. This happens transparently for the WWW server. The applet notifies a tracking component on the server about its start up and shutdown, which correspond to the beginning and end of the period the document is displayed. The tracking system is just one source for spatial information about users, other sources include the WWW server log files and proxies.

Work on developing a similar system for audio is currently under way.

III.A.2 Web only Approach

The second approach evaluated was to use the Common Gateway Interface (CGI) support provided by most servers. The CGI allows servers to delegate responsibility for generating a web page to an external program. This interface is generally used to handle user input through form based questionnaires. It is now also common for this approach to be used by sites with databases, a web page can be created where the content is filled in by a program interrogating the database in response to a users request.

If web pages are generated on-the-fly their content does not have to be pre-defined, in other words the same request sent by different users can result in two entirely different pages. This is useful because pages can then be made to contain the current state of the system. The importance of pages being able to carry around state information becomes apparent when one considers that web servers are

effectively stateless, and retain no information about page requests. This idea was used at Lancaster a few years ago to construct a virtual map of the department, through which people could wander and meet members of the department, and is now used in a diverse range of applications including the common search engines.

This approach has allowed a CoBrow to be built with dynamic creation and updates of meeting places and conferences within meeting places without daemons running on client machines, in fact the system works with no server or client modifications.

An interesting side effect of using the CGI approach is that the system is proxy safe. One problem a CoBrow system faces is that proxy servers often respond to user requests by sending previously cached copies of pages. Unfortunately this means the server that originally contained the page never sees the request and therefore it is difficult to maintain a central record of the people asking for the page. A CGI based system can get round this by requesting that all requests for the base text of a page be referred back to the main server on which the page resides. Proxy servers will then only respond to requests for objects within the page, for example, images. This means we get the traffic reducing advantages of using proxy servers but still have a central record of requests for each page.

Although shared browsing capabilities are provided within the user's browser, multimedia conferencing functionality is provided by external applications which can be downloaded and registered by the user. Once registered as a helper application, the CoBrow system can invoke all the tools necessary for participation in a conference automatically.

The conferencing tools available are not restricted by the system as it was felt important that new tools should be available when released, and that users should be able to choose for themselves the best set of tools for the type of conference they wish to have. Currently the recommended tool set includes CUSeeMe, and the Internet Multicast Backbone (MBone) tools. Management of created conferences is handles by the CoBrow system and no user configuration of the tools is necessary.

Within this prototype we have also experimented with the use of multiple servers. If successful, there would be many CoBrow servers around the Internet, each managing meeting places and conferences for a relatively small number of users - corresponding to groups or departments. For the system to work effectively, servers must be able to share information on the meeting places, conferences and users they manage. This system handles this by providing a CGI program that generates a complete dump of a server's local state, remote servers can pull this information to complete their view of global state.

A drawback of this approach is that the user interface is limited by the available HTML tags, which currently only allow simple buttons, menus, and text entry boxes.

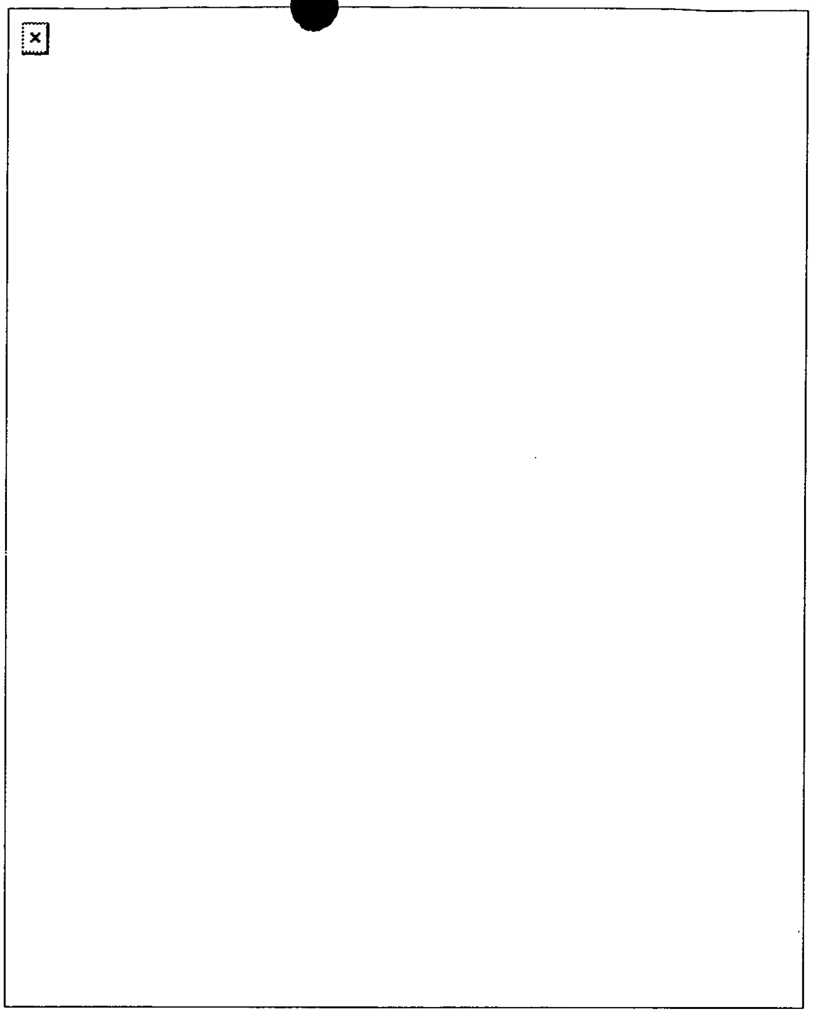

Figure 4: A browser showing the Web-only interface

III.A.3 Java Approach

A very interesting development for designers of web based applications is the recent addition of Java support within the main web browsers. Java allows complex applications to be developed in a proper object oriented language; these applications can then form part of a conventional web page and appear within the displayed text in the same way as images, etc.

The use of Java means that complex user interfaces, far superior to those previously possible within web pages, can be built quite easily. The level of sophistication possible with this approach can be seen in a number of recently developed systems, for example, the Java implementation of the Corel Office suite [2]. There is considerable experience with Java within the CoBrow team from other projects and all the partners are keen to see Java based interfaces adopted in the final implementation of CoBrow.

As Java is an interpreted language, run within the web browser, user interfaces can be far more responsive than possible using the simple *submit* and *timed-response* techniques available to CGI based systems. This opens up a whole new class of more interactive interfaces.

The first interface to be tried was a simple text based chat application where users could discuss material using similar functionality to that available through applications such as Unix *talk* and *irc*. The more immediate response and feedback from the system results in a very useable interface, however,

due to the interpreted nature of Java and the resulting low performance there is still a place for more conventional collaborative tools, such as the MBone software mentioned earlier.

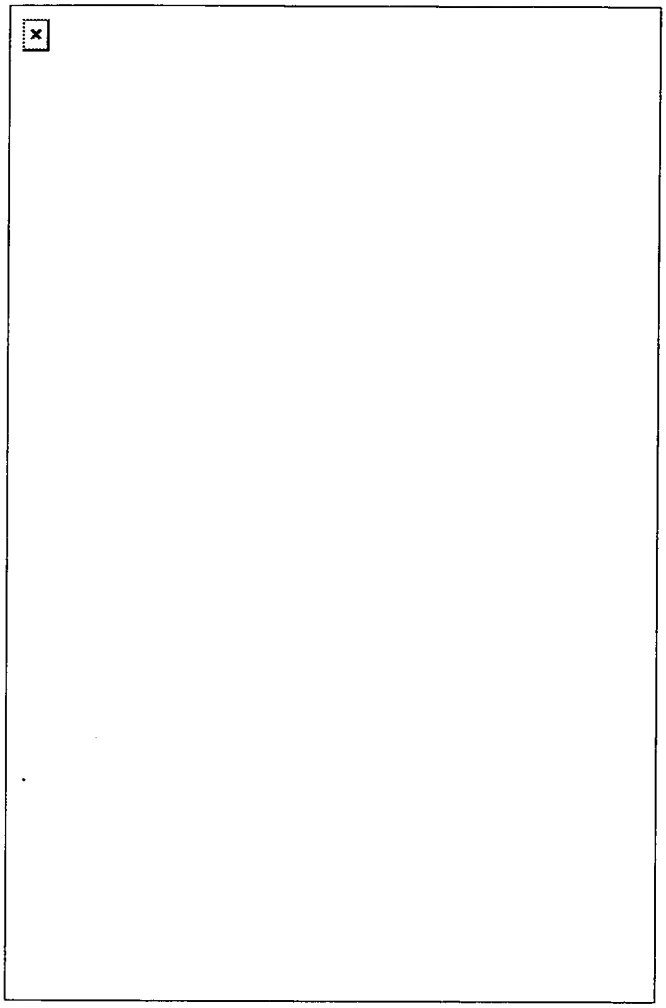

Figure 5: Sophisticated user interfaces are possible with Java

IV. Conclusions

The development of prototypes showing different types of user interfaces and also different approaches to collecting and collating the information needed to maintain and operate a CoBrow system has proved extremely valuable. We have gained considerable knowledge of the relative performance and efficiency of the various underlying implementation techniques and built three systems able to support user-trials.

Our experiences with Java have removed any doubt in our minds that a collaborative system built around the web must be based on Java like technology, and as we move from the prototype stages of the project to the development of the final system we see more and more emphasis on Java.

Early feedback from users, following small scale trials of the three systems, has been positive. Users have found the ability to contact other users with similar interests without prior arrangement and without previous knowledge of their existence a useful addition to the familiar web environment.

It has also been interesting to note the almost immediate adoption, of even the early and more restricted versions, of the prototypes among the developers for their day to day work. We believe this and the level positive user feedback clearly shows a need for such collaborative systems.

V. Ongoing and Future Work

An issue that must be addressed is how a multi-cast based service could be efficiently extended to users connecting via Internet providers over modem links. An even greater problem is the growing demand for mobility to be considered in the development of this type of service - IPv6 [4] should help but only by allowing movable endpoints to connections, the problem of variable levels of service and connectivity will still be a significant issue.

The wide scale adoption of proxy servers now prevents many page requests getting through to the main server for the page, instead an intermediate proxy server will reply with a previously cached copy of the page. Getting an accurate picture of which users are looking at which pages is obviously an important requirement when trying to establish vicinities, and unfortunately obtaining such information is made far more difficulty with the advent of proxy servers. Any real collaborative browsing system must work correctly even when users have proxy servers configured. The 'web only' prototype avoids this problem by making use of a characteristic of the page retrieval mechanism, but this can not easily be transferred to either of the other prototypes - more work clearly has to be done in this area.

Another issue for 'real' systems is support for systems behind firewalls. Firewalls can place severe restrictions on connections to hosts outside an institution and could easily reduce the effectiveness of a collaborative conferencing system. There is a significant, and growing, number of particularly commercial organisations with installed firewalls and so solutions must be found to these problems without compromising the security which these firewalls have been installed to provide.

Finally, in the longer term, more sophisticated ways of identifying users' interests need to be investigated, and also ways in which these more complex definitions of interest can be matched in order to support and enable collaboration. There is obviously scope for investigating existing work on automatic abstract generation.

VI. Acknowledgments

This work has been performed in the framework of the project CoBrow, which is partly funded by the European Community under the Telematics for Research program and the Swiss Department of Education and Science. The authors would also like to thank their colleagues on the CoBrow project at ETH Zurich, Lancaster University, UK, and at the University of Ulm, Germany for their contributions.

VII. References

[1] AltaVista Search, Nov. 1996.

[2] Corel Corporation homepage, Nov. 1996.

[3] Forrester Research, "Forecast of users of on-line services", Computer Zeitung, Vol. 8, 23. Feb.

1995.

[4] C. Huitema, "IPv6 - The New Internet Protocol", Prentice Hall, ISBN 0-13-241936-X, 1996.

[5] Interactive Model Railroad, Universität Ulm, Nov. 1996.

[6] The CoBrow Project, Nov. 1996.

Author Information

Dr. Andrew Scott is a lecturer at the computing department of Lancaster University. His research interests include computer supported collaborative work, multimedia, high speed networking and wireless networking.

Gabriel Sidler received a diploma in computer science from HTL Brugg-Windisch, Switzerland and a master's degree in electrical engineering from Northwestern University, Chicago. He is currently a research assistant at the computer engineering and networking laboratory of the Swiss Federal Institute of Technology, Zurich. His research interests include multimedia communication and computer supported collaborative work.

Heiner Wolf is a research assistant at the distributed systems department of the University of Ulm, where he is pursuing a Ph.D. degree in computer science. His research interests include computer supported collaborative work, multimedia, high speed networks and ATM.

Security and Privacy Aspects in a Collaborative Browsing Environment

Telematics for Research, Project RE1003, CoBrow Collaborative Browsing in Information Resources

Deliverable D4.8

Version 1.1

Author: Marcel Dasen

Contributions: Germano Caronni, Gabriel Sidler

Summary

The deliverable covers to main topics:

- How CoBrow can be run across firewalls
- How security and privacy can be achieved in CoBrow.

Computer systems connected to public networks, such as the Internet, are generally endangered of being compromised in their integrity by intruders gaining unauthorized access to resources not intended by the owner of the computer. Flaws in design and implementation of operating systems and communication software, which allow for such access are almost always only discovered after a successful attack. To prevent damage to the vital information resources, many sites have installed *firewall-systems*, which act as gateways between the internal - and the public networks. Firewalls are indispensable for securing internal, private networks connected to public networks. For CoBrow firewalls turn out to be a major obstacle, which can only be dealt with by making them an integral part of design. This document shows how CoBrow can be made as compatible as possible with existing firewalls, but also shows the limitations imposed by certain security policies of firewalls. It is concluded that firewalls turn out to be a major obstacle and certain security precautions will prevent CoBrow from running across the firewall at some sites:

- Sites that have HTTP disabled
- Sites that filter IP packets from known CoBrow servers
- Sites that neither allow multicast traffic over the MBone nor allow a user's host to become a server for multimedia traffic, will not be able to do synchronous multimedia communication such as audio/video conferencing.

CoBrow is a distributed system used for collaboration and communication of WWW-users. Such a system is required to offer a reasonable amount of privacy to its users and, therefore, CoBrow has to be extended by security mechanisms. Meeting other people through CoBrow requires the users to provide some data describing himself. For users to be willing to give up on some anonymity they have to trust the system to only use this data as advertised. Since the components can be runned by everyone an authentication scheme is needed for the components, which ensures that only trustworthy organisations can take part in this information exchange process. This paper proposes to use a registration process for CoBrow servers to be able to tell the user who might get the provided information and to identify the organisation who misused the system. Furthermore, it is analysed how users can be identified when entering the system. Two protocols are proposed: A login/password scheme for applications where only little authenticity is needed and public-key cryptography if a higher degree of confidence in the identification is required. It is explained how this can be accomplished by applying standard encryption technology (PGP) and security protocols (Secure-HTTP). Ease of use is considered very important. Thus the chosen security technology is carefully scaled to the risks involved. The Diffie-Hellman protocol is proposed for setting up secure channels for close user groups and private information exchange. This famous protocol's patent just expired end of April this year and thus Diffie-Hellman can be freely applied.

---

yourEmail@xyzcorp.com

Collaborative Browsing in the World Wide Web (1997) (Make Corrections) (19 citations)
Gabriel Sidler, Adrew Scott, Heiner Wolf

View or download:
terena.nl/conf/jenc8/papers/122.ps
Cached: PS.gz PS PDF DjVu Image Update Help

From: terena.nl/conf/jenc...proceedings (more)
Homepages: G.Sidler A.Scott [2] [3] H.Wolf [2] [3] HPSearch (Update Links)

CiteSeer Home/Search Bookmark Context Related

(Enter summary)

Rate this article: (best)
Comment on this article

Abstract: The World Wide Web (WWW) is today the most successful service of the Internet. The richness of information available combined with easy access to this information makes it a premier information gathering tool for researchers and consumers, however, the model of today's WWW does not include the users. The WWW is a purely information focused environment, consisting of documents and links between these documents. The virtual world formed by the linked information on the WWW is completely separated ... (Update)

Context of citations to this paper: More

.... **space, such as providing users with awareness of other users currently viewing the same web page, as demonstrated in the CoBrow system [11].** The work of Benford et al. on mixed reality boundaries [2] has also inspired much of our current work and is concerned with the merging...

...same place at the same time. **In this context, awareness support is provided to facilitate social communication and collaborative browsing [14].** An obvious example for facilitation of communication through web places are chat rooms. In more elaborate systems, audio video...

Cited by: More
Virtual Reviewers for Collaborative Exploration - Of Movie Review (Correct)
Sharing (Location) Context to Facilitate.. - Cheverst.. (2000) (Correct)
Easy Teach & Learn: A Web-Based Adaptive Middleware for.. - Walter, Ruf, Plattner (2000) (Correct)

Similar documents based on text: More All
0.2: Database and Modelling Strategy: A Compliant Way for.. - Nathalie Farenc Frederic (2000) (Correct)
0.1: Mobile Wideband Terminal Applications - The Magic (Correct)
0.1: The Multimedia Session Manager User Manual - Rafaeli, Hutchison (Correct)

Related documents from co-citation: More All
5: Designing Calm Technology (context) - Weiser, Brown - 1996
4: Challenges and Perspectives for Cooperative Work on the Web (context) - Dix - 1996
4: GroupWeb: A WWW browser as real time groupware (context) - Greenberg, Roseman - 1996

BibTeX entry: (Update)

G. Sidler, A. Scott, and H. Wolf. Collaborative Browsing in the World Wide Web. In Proceedings of the 8th Joint European Networking Conference, Edinburgh, Scotland, May 1997. http://citeseer.nj.nec.com/sidler97collaborative.html More

```
@misc{ sidler97collaborative,
  author = "G. Sidler and A. Scott and H. Wolf",
  title = "Collaborative Browsing in the World Wide Web",
  text = "G. Sidler, A. Scott, and H. Wolf. Collaborative Browsing in the World Wide
    Web. In Proceedings of the 8th Joint European Networking Conference, Edinburgh,
    Scotland, May 1997.",
  year = "1997",
  url = "citeseer.nj.nec.com/sidler97collaborative.html" }
```

Citations (may not include all citations):
1 Forecast of users of on-line services (context) - Research - 1995
1 uni-ulm (context) - Railroad, rrvs - 1996

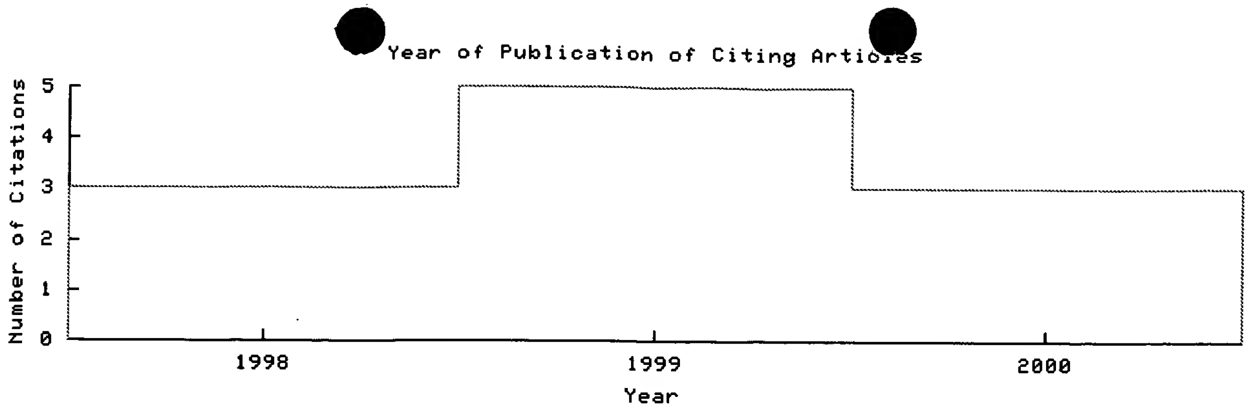

The graph only includes citing articles where the year of publication is known.

Documents on the same site (http://www.terena.nl/conf/jenc8/proceedings.html): More
Recent Activities in the MERCI Conferencing Project - Kirstein, Bennett (Correct)
NATO's contribution to Computer Networking in Central and.. - Nadreau Nadreau (Correct)
Desktop Video: Building Large-scale Services with.. - Bogen, Bonkowski, .. (Correct)

Reading: Collaborative Browsing

Username: maccara
Last name: MacCara
First name: Kori
Course: CSCI 6504
Due: June 15, 1999

Papers Reviewed:

Authors: Lieberman, Van Dyke, and Vivacqua
Title: Let's Browse: A Collaborative Web Browsing Agent

Report

This paper focused on the many problems associated with group browsing. According to the authors, group browsing is quite common and should be addressed. There are many problems that are incurred when attempting to accommodate the common interests of the group when browsing. Such problems include acknowledging the interests of each person in the group, and doing so in a timely fashion. The authors also state that groups are in general more impatient than individuals, therefore the timely training processes of individual Web agents could not be used. Another problem associated with group browsing is detecting when an person joined or left the group, accommodating for such changes within a group must be accounted for.

The solutions to these problems presented in this paper involved first a way to identify each individual to the browser agent. In this case, active tags were used to transmit a users identity to the browser using infrared transfer (as long as the person was inside the transmit radius and in line-of-sight of the infrared sensor of the browser). Once the identity of a person is known the agent must acquire some knowledge about this persons interests to incorporate into the group information. To gather this information, the agent is given a time limit to scan the persons Web page or their business page if a personal page is not available. Using TFIDF term frequency measurements, approximately 50 terms that are significant to the individual are retrieved from their pages. These terms can then be used to find interest overlaps with the other members of the group. When a person leaves or joins the group the agent must re-evaluate the group browsing preferences and begin anew to generate the list of pages that are of potential interest to the group.

According to the authors of this paper, this solution received favorable reviews from users, although no formal controlled tests were performed. This solution was able to generate a list of terms that were of interest to a user in a set time frame, which compromised the accuracy of the users preferences in the interests of time. This reduced the amount of frustration the group would experience due to processing time and produced reasonably accurate user profiles.

The large number of terms retrieved proved to be more useful in a group setting as opposed to the small number of terms in an individual profile. This is because terms that are not very significant for an individual can increase in weight if it overlaps with the profile of another member of the group. Through the use of a larger term profile for a user, there is a higher chance that there will be a common overlap with other members of the group.

Although the system seemed to be generally well accepted by users, the authors proposed a series of improvements that could not only increase the usability of the system, but also the accuracy of the group profile by generating better individual profiles. Some suggestions included the use of multiple pointing devices, which can be associated with a user. Another suggestion would be to alternate the weights associated with each user profile so that no one user profile can dominate the group. Such interactive input and weighting schemes can only help to further train the system and create an evolving group profile which will better represent the interests of each member.

Collaborative Bro[illegible]g in the World Wide Web
(1997) (Make Corrections) (19 citations)
Gabriel Sidler, Adrew Scott, Heiner Wolf

View or download:
terena.nl/conf/jenc8/papers/122.ps
Cached: PS.gz PS PDF DjVu Image Update Help

From: terena.nl/conf/jenc...proceedings (more)
Homepages: G.Sidler A.Scott [2] [3]
H.Wolf [2] [3] HPSearch (Update Links)

CiteSeer Home/Search Bookmark Context Related

(Enter summary)

Rate this article: (best)
Comment on this article

Abstract: The World Wide Web (WWW) is today the most successful service of the Internet. The richness of information available combined with easy access to this information makes it a premier information gathering tool for researchers and consumers, however, the model of today's WWW does not include the users. The WWW is a purely information focused environment, consisting of documents and links between these documents. The virtual world formed by the linked information on the WWW is completely separated ... (Update)

Context of citations to this paper: More

.... **space, such as providing users with awareness of other users currently viewing the same web page, as demonstrated in the CoBrow system [11].** The work of Benford et al. on mixed reality boundaries [2] has also inspired much of our current work and is concerned with the merging...

...same place at the same time. **In this context, awareness support is provided to facilitate social communication and collaborative browsing [14].** An obvious example for facilitation of communication through web places are chat rooms. In more elaborate systems, audio video...

Cited by: More
Virtual Reviewers for Collaborative Exploration - Of Movie Review (Correct)
Sharing (Location) Context to Facilitate.. - Cheverst.. (2000) (Correct)
Easy Teach & Learn: A Web-Based Adaptive Middleware for.. - Walter, Ruf, Plattner (2000) (Correct)

Similar documents based on text: More All
0.2: Database and Modelling Strategy: A Compliant Way for.. - Nathalie Farenc Frederic (2000) (Correct)
0.1: Mobile Wideband Terminal Applications - The Magic (Correct)
0.1: The Multimedia Session Manager User Manual - Rafaeli, Hutchison (Correct)

Related documents from co-citation: More All
5: Designing Calm Technology (context) - Weiser, Brown - 1996
4: Challenges and Perspectives for Cooperative Work on the Web (context) - Dix - 1996
4: GroupWeb: A WWW browser as real time groupware (context) - Greenberg, Roseman - 1996

BibTeX entry: (Update)

G. Sidler, A. Scott, and H. Wolf. Collaborative Browsing in the World Wide Web. In Proceedings of the 8th Joint European Networking Conference, Edinburgh, Scotland, May 1997. http://citeseer.nj.nec.com/sidler97collaborative.html More

```
@misc{ sidler97collaborative,
  author = "G. Sidler and A. Scott and H. Wolf",
  title = "Collaborative Browsing in the World Wide Web",
  text = "G. Sidler, A. Scott, and H. Wolf. Collaborative Browsing in the World Wide
    Web. In Proceedings of the 8th Joint European Networking Conference, Edinburgh,
    Scotland, May 1997.",
  year = "1997",
  url = "citeseer.nj.nec.com/sidler97collaborative.html" }
```

Citations (may not include all citations):
1 Forecast of users of on-line services (context) - Research - 1995
1 uni-ulm (context) - Railroad, rrvs - 1996

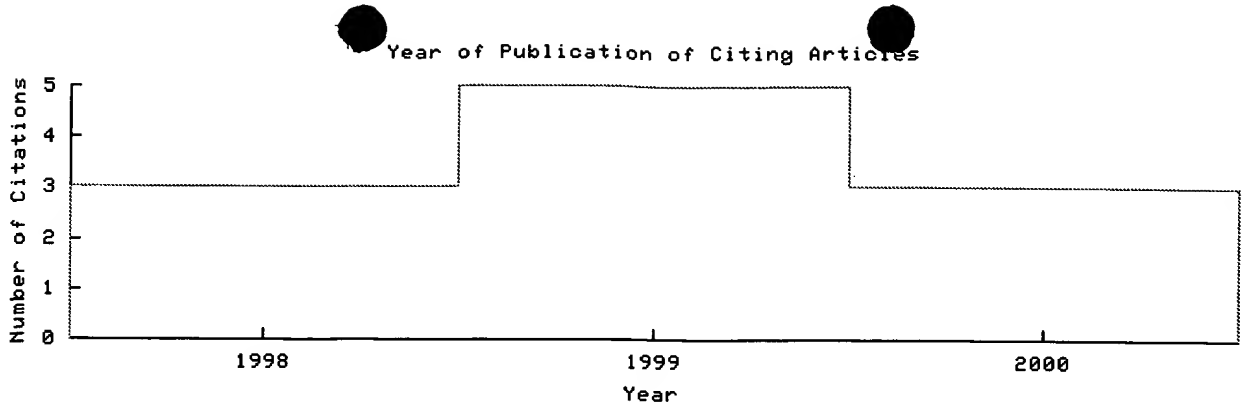

The graph only includes citing articles where the year of publication is known.

Documents on the same site (http://www.terena.nl/conf/jenc8/proceedings.html): More
Recent Activities in the MERCI Conferencing Project - Kirstein, Bennett (Correct)
NATO's contribution to Computer Networking in Central and.. - Nadreau Nadreau (Correct)
Desktop Video: Building Large-scale Services with.. - Bogen, Bonkowski, .. (Correct)